

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 917 044 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
19.05.1999 Bulletin 1999/20

(51) Int. Cl.⁶: G06F 3/12

(21) Application number: 98309145.5

(22) Date of filing: 06.11.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

- Fan, Sheng Kuo
Los Angeles, California 90064 (US)
- Fung, Joseph Z
Artesia, California 90701 (US)
- Nesbitt, David P
Redondo Beach, California 90278 (US)

(30) Priority: 07.11.1997 US 966409

(71) Applicant: XEROX CORPORATION
Rochester, New York 14644 (US)

(74) Representative: Skone James, Robert Edmund
GILL JENNINGS & EVERY
Broadgate House
7 Eldon Street
London EC2M 7LH (GB)

(72) Inventors:
• Tan, Shee-Yen
Walnut, California 91789 (US)

(54) Improved load balancing of distributed printing systems using enhanced printer attributes

(57) A distributed printing system that more optimally distributes job requests within a distributed printing system based on more detailed information about the particular attributes of each output device. In particular, additional printer attributes are used that allow greater utilization of a plurality of attached printers having differing capabilities. These attributes include "max-concurrent-jobs-processing" (MCJP), "number-of-jobs-on-

device" (NJOD) and "number-of-cascaded-jobs" (NCJ), and a attribute value "saturated" for the attribute "printer-state". The attribute MCJP is an integer value that is based on the capabilities of the printer. NJOB and NCJ attributes allow monitoring of jobs sent to each printer for supported and unsupported printers. If NJOB or NCJ are greater than MCJP, then the attribute "printer-state" is set to "saturated", prohibiting further spooling of jobs to that particular printer.




FIGURE 1

BEST AVAILABLE COPY

EP 0 917 044 A2

user file that is to be printed. A print job represents a collection of one or more documents that are printed as a unit. The print job includes instructions for printing (such as production and finishing), event notification, and response delivery. The server 400 also is responsible for sending job data and associated job control commands to the output device 500.

[0018] Each output device 500 is a physical device or hardware that is capable of rendering images or documents to produce a hard copy output of the print jobs received from the server 400. Depending on the type of output device 500 and/or the settings of its particular features, each output device 500 returns responses to the server 200 indicating, for example, its current state. Output devices 500 can include, but are not limited to, printers, facsimile machines or pagers. However, as a nonlimiting example, this disclosure refers to the output devices 500 as printers, for simplicity.

[0019] In a preferred system, implemented on a general purpose computer, the server functions are split between two different server types, a spooler server 410 and a supervisor server 420. Additionally, such an exemplary system supports two types of output devices: logical printers 412 (LPI, LP2 and LP3) and physical printers 422 (PP1 and PP2).

[0020] The spooler 410 receives a client print request for a print job that specifies a particular logical printer. The spooler 410 then schedules the print job on an appropriate physical printer 422 associated with the specified logical printer 412. At an appropriate time, the spooler 410 forwards the print job to the supervisor 420 associated with the physical printer 422 on which the spooler 420 scheduled the print job. The spooler 410 can support multiple supervisors 420.

[0021] The supervisor server 420 delivers data to one or more of the output devices 500. The supervisor 420 also receives jobs from the spooler 410, interprets job requests for print submission, receives print data, passes the data and the job control commands to an appropriate output device 500, and handles any responses made by the output devices 500.

[0022] The supervisor 420 receives client management requests that apply to print jobs forwarded to the supervisor, including those print jobs forwarded further to one of the output devices 500. The supervisor 420 also sends management requests to the spooler 410, such as, for example, a request for next job, and status updates, such as "job completed." The supervisor 420 also notifies the spooler 410 of events, such as canceling a print job in response to a "Job Cancel" management request from a user.

[0023] The logical printer 412 indicates particular characteristics and capabilities of one or more physical printers 422. In essence, a logical printer is an abstraction of printer capability that serves as a selectable device and a gateway into the spooler 410. A physical printer 422 represents an actual output device, such as a Xerox 4230 laser printer connected to the printing sys-

tem 100.

[0024] Users of the client 200 direct print requests to logical printers 412 having characteristics that meet their needs, such as A4 paper, highlight color, and/or stapled output. The spooler 410 channels the request via the queue 414 to the supervisor 420 that supports a physical printer 422 mapped to the specified logical printer 412 through the queue 414. The queue 414 associates a set of print jobs with a set of logical and physical printers 412, 422 and hold print jobs until the spooler 410 can send the print jobs to an appropriate supervisor 420. Logical printers 412 feed print jobs into the queue 414 and the physical printers 422 request print jobs from the queue.

[0025] As shown in Figure 2, a larger network client/server system can have numerous logical printers associated with numerous physical printers through a plurality of queues. This type of system is capable of load balancing, in which jobs can be directed to alternate printers associated with the same queue. The alternate printers also provide printing characteristics required by an associated logical printer.

[0026] Such a distributed printing system uses objects to represent various entities, such as the logical printers, the physical printers and/or the queues, although objects can represent the spoolers and the supervisors as well. An object contains a collection of attributes that provide information about the object. Examples of printer attributes are:

printer-name: a single-valued attribute whose value is a name that uniquely identifies a logical or physical printer.

media-supported: a multi-valued attribute whose values identify the media that the printer supports, such as "iso-A4 white".

printer-state: a single-valued attribute that indicates the current status of a printer, such as "idle" or "printing".

[0027] Each object is defined by this set of attributes, which can contain static and dynamic attributes. The spooler 410 and supervisor 420 manipulate the object representation of an entity they are operating on by setting or changing the values of the object's attributes. For those attributes that can be manually changed, a user can observe the value of an attribute with a "list attributes" command and change the value of an attribute with a "set attributes" command. The print capabilities of a print system are limited by the predefined set of system attributes fixed for each printing system.

[0028] Figure 3 shows the interaction of various components of a preferred client/server printing system when a print job is submitted for processing. This system includes a supervisor server 420 that includes a front end module 424, a document processor 426, and a back end module 428. There can be multiple back end modules 428 complemented in the supervisor 420, as de-

performed based on the number of concurrent processing jobs the particular output device can handle. The invention recognizes that various printers associated with the printing system can have greatly varying capabilities in speed and other qualities. For instance, the distributed system may include a desktop printer with a very low output speed of around 3-4 ppm and an production printer, such as the Xerox DocuPrint 4635 NPS, which has a very high output speed of about 135 ppm. Printers also have varying levels of internal memory for handling and queuing jobs. Accordingly, by this invention, individual printers are assigned a MCJP attribute based on the number of jobs that can be effectively handled concurrently.

[0044] Taking a production printer having a MCJP attribute set to 5 as a nonlimiting example, the document processor 426 will not stop sending out print job requests to the spooler 410 until the query status of the output device 500 shows five queued jobs to the output device 500. While the number of jobs is less than the maximum for the output device 500, the document processor 426 will inform the front end module 424 to periodically request for job in order to balance the load.

[0045] If the "number-of-jobs-on-device" or number-of-cascaded-jobs value of these attributes exceeds the set value "max-concurrent-job-processing" attribute, the "printer-state" attribute is set to "saturated" and jobs are no longer submitted to the associated remote spooling server for that particular printer. This feature balances the load of system printers, no matter if the printers [probably not necessary as we are only concerned with the ability to return a NJOD value.] running on a foreign printing system could return the NJOD or not.

[0046] A specific example of such a system is illustrated in Fig. 4 in which supervisor 420 manages four physical printers 502, 504, 506 and 508, each of varying capabilities. Printer 502 is a low-end desktop printer. Printers 504 and 508 are production printers with higher capacities and much improved memory and processing speed. Printer 506 is a production printer of even greater speed and capability, such as the Xerox DocuPrint 4635 printer.

[0047] The attribute database 600 associated with the supervisor 420 includes attributes of each printer, of which only a partial attribute list is shown. Individual MCJP values are entered by a system administrator based on the individual capabilities of each particular printer. These values are integer values that represent the number of concurrent jobs that can be handled by the device. Exemplary values are "1" for the desktop printer, "5" for the production printer, and "10" for the high speed production printer. However, these values can be arbitrarily changed to more accurately reflect the particular capabilities of the output device relative to the particular distributed printing system used and other output devices on the printing system. The other listed attributes are dynamic and change upon changes in the system load.

[0048] In this example desktop printer 502 is shown in an "idle" printer state with no jobs cascaded to or on the device. The printer 502 is obviously ready to accept a print job. The production printer 504 is in a "printing" printer state. As the attribute NJOD is NULL, this printer does not support this function. However, job tracking can still be achieved by use of the NCJ attribute, which indicates that 3 jobs have been cascaded (sent) to the device. As this particular printer, a production printer, is capable of handling up to five concurrent jobs as indicated by the MCJP=5 attribute, printer 504 is capable of receiving additional jobs (i.e., the attribute printer-state is not set to "saturated").

[0049] Likewise, the printer 506, a high speed production printer is currently in a "printing" printer state and is able to accept additional print jobs, as it currently has less than the maximum number of print jobs it can handle. The printer 506 supports the NJOD feature. As this attribute is more accurate than NCJ, it is used in determining load balancing. However, the printer 508, another production printer, is not able to accept any more jobs as its printer state is "saturated" due to the NJOD (as well as NCJ) value being equal to or greater than the MCJP value of 5. Accordingly, job requests are no longer sent for PP4 until its status returns to a "printing state" and the attribute printer-state has been set to "saturated" (i.e., when the number of jobs on device becomes less than MCJP).

[0050] Fig. 5 is a flowchart outlining one method for processing print requests. In step S510, the control routine initializes the attributes file for each printer corresponding to the attributes file 600. As the value for MCJP is a static value and is preset by the system administrator using a common input, such as a Graphical User Interface (GUI), during programming or setting up of the corresponding output device 500, its value remains the same until the system administrator desires a change in this value. The attribute NCJ is set to "0" for each printer, as no jobs are pending. The attribute NJOD is set by default to "NULL", but may be set to zero if this feature is supported by the corresponding printer. The printer-state (PS) of each printer should be "idle".

[0051] In step S520, if the printer state is saturated, the control routine then conducts a status check for the "printer-state". If the printer is in idle or printing state, control continues to step S530 where the print request is sent to spooler 410 and queue 414 is asked to schedule a job. If, in step S540, there is no new job available, control jumps back to step S540. Once there is a new job available to be scheduled by queue 414, which is notified by logical printer 412, control continues to step S550. Then, in step S560, the print job is sent to a non-saturated physical printer 422 and the attribute NCJ is incremented by 1. Then, in step S570, the supervisor 420 conducts print job status queries with the physical printer 422 to determine the status of the print job.

[0052] Fig. 6 shows a more detailed printing process. In step S605, the printer attributes are initialized and in

5. A system according to any of the preceding claims, wherein the printing system includes an updating means for updating the "number-of-jobs-on-device" attribute in the attribute database based on a query response received from the at least one output device.

6. A method of performing load balancing in a distributed printing system having at least one output device, comprising:

storing printer attributes in a database, the printer attributes including:

an attribute "max-concurrent-jobs-processing" assigned a numeric value corresponding to a number of concurrent print jobs that can be handled by the at least one output device,
an attribute "printer-state" having a value "saturated", and at least one of:

an attribute "number-of-cascaded-jobs" that is a dynamic attribute having a value that is incremented by the printing system upon sending a print job to the at least one output device,
and
an attribute "number-of-jobs-on-device" that is a dynamic attribute that is updated by the at least one output device;

receiving a print request;
generating a print job from the print request and storing the print job until the print job can be sent to the at least one output device;
comparing at least one of the attributes "number-of-jobs-on-device" and "number-of-cascaded-jobs" attributes against the attribute "max-concurrent-jobs-processing" and setting the attribute "printer-state" to "saturated" for the at least one output device when either has a value greater than or equal to the value of the attribute "max-concurrent-jobs-processing";
and
sending the print job to the at least one output device if the attribute "printer-state" is idle or printing.

7. The method of claim 6, wherein the printing system is a client/server distributed printing system comprising a client, a spooler, a supervisor, and at least one output device, the method further comprising:

receiving a print request from the client at the spooler;
generating a print job from the request and stor-

ing the print job in a queue within the spooler;
generating a job request at the supervisor when at least one of the output devices has the attribute "printer-state" value equal to idle or printing; and
sending the print job to the supervisor in response to the job request for subsequent sending to the at least one output device.

8. A method according to claim 6 or claim 7, further comprising incrementing the attribute "number-of-cascaded-jobs" upon sending the print job to the at least one output device.

9. The method of claim 8, further comprising decrementing the attribute "number-of-cascaded-jobs" upon indication of completion of the print job by the at least one output device.

10. A method according to any of claims 6 to 9, wherein at least two output devices are provided and respective "max-concurrent-jobs-processing" attributes are different based on differing capabilities of the respective output devices.

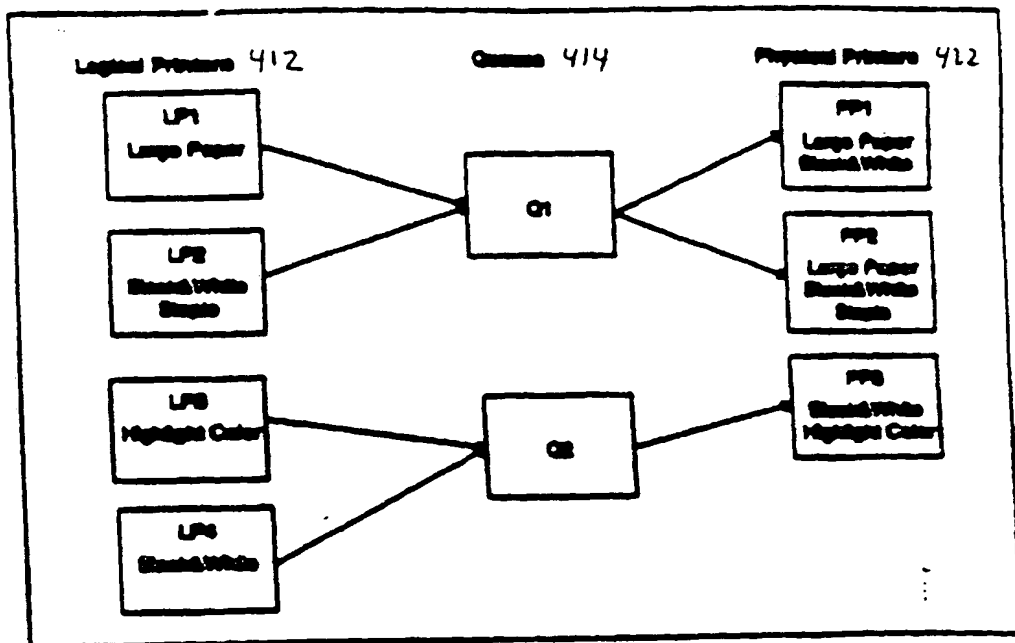



FIG. 2



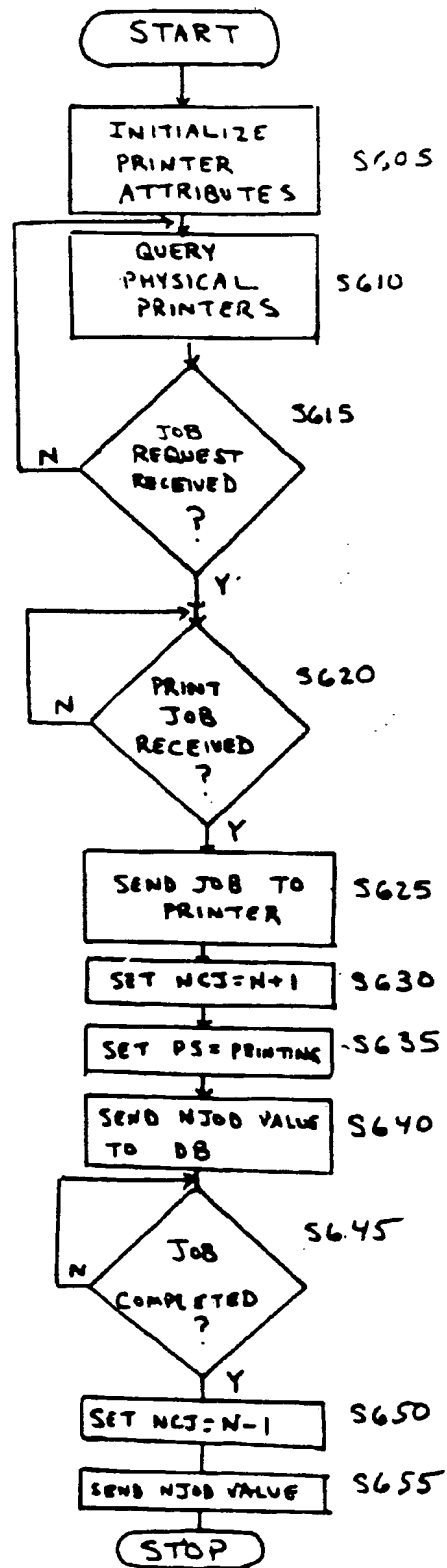



FIG. 6



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 917 044 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
03.05.2000 Bulletin 2000/18

(51) Int Cl.7: G06F 3/12

(43) Date of publication A2:
19.05.1999 Bulletin 1999/20

(21) Application number: 98309145.5

(22) Date of filing: 06.11.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

- Fan, Sheng Kuo
Los Angeles, California 90064 (US)
- Fung, Joseph Z
Artesia, California 90701 (US)
- Nesbitt, David P
Redondo Beach, California 90278 (US)

(30) Priority: 07.11.1997 US 966409

(71) Applicant: XEROX CORPORATION
Rochester, New York 14644 (US)

(74) Representative: Skone James, Robert Edmund
GILL JENNINGS & EVERY
Broadgate House
7 Eldon Street
London EC2M 7LH (GB)

(72) Inventors:
• Tan, Shee-Yen
Walnut, California 91789 (US)

(54) Improved load balancing of distributed printing systems using enhanced printer attributes

(57) A distributed printing system that more optimally distributes job requests within a distributed printing system based on more detailed information about the particular attributes of each output device. In particular, additional printer attributes are used that allow greater utilization of a plurality of attached printers having differing capabilities. These attributes include "max-concurrent-jobs-processing" (MCJP), "number-of-jobs-on-

device" (NJOD) and "number-of-cascaded-jobs" (NCJ), and a attribute value "saturated" for the attribute "printer-state". The attribute MCJP is an integer value that is based on the capabilities of the printer. NJOD and NCJ attributes allow monitoring of jobs sent to each printer for supported and unsupported printers. If NJOD or NCJ are greater than MCJP, then the attribute "printer-state" is set to "saturated", prohibiting further spooling of print jobs to that particular printer.




FIGURE 1

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 98 30 9145

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

13-03-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5559933 A	24-09-1996	NONE	
EP 0697648 A	21-02-1996	US 5715381 A	03-02-1998
		JP 8063341 A	08-03-1996

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82